

A discrete approach to stochastic parametrization and dimensional reduction in nonlinear dynamics

Alexandre J. Chorin and Fei Lu

Department of Mathematics, University of California, Berkeley, CA 94720, USA.
Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA.

Abstract Many physical systems are described by nonlinear differential equations that are too complicated to solve in full. A natural way to proceed is to divide the variables into those that are of direct interest and those that are not, formulate solvable approximate equations for the variables of greater interest, and use data and statistical methods to account for the impact of the other variables. In the present paper the problem is considered in a fully discrete-time setting, which simplifies both the analysis of the data and the numerical algorithms. The resulting time series are identified by a NARMAX (nonlinear autoregression moving average with exogenous input) representation familiar from engineering practice. The connections with the Mori-Zwanzig formalism of statistical physics are discussed, as well as an application to the Lorenz 96 system.

1 Introduction and outline

There are many problems in science where the equations of motion are too complex for full solution, either because the equations are not certain or because the computational cost is too high, but one is interested only in the dynamics of a subset of the variables. Such problems appear, for example, in weather and climate modeling, see e.g. [1, 2], in economics, see e.g. [3], in statistical mechanics, see e.g. [4, 5], and in the mechanics of turbulent flow, see e.g. [6, 7]. In these settings it is natural to look for simpler equations that involve only the variables of interest, and then use data to assess the effect of the remaining variables on the variables of interest in some approximate way. In the present paper we focus on stochastic methods for doing this, with data coming either from experimental observations or from prior numerical calculations.

Consider a set of differential equations of the form:

$$\frac{d}{dt}x = R(x, y), \quad \frac{d}{dt}y = S(x, y), \quad (1)$$

where t is the time, $x = (x_1, x_2, \dots, x_m)$ is the vector of resolved variables, and $y = (y_1, y_2, \dots, y_\ell)$ is the vector of unresolved variables, with initial data $x(0) = \alpha$, $y(0) = \beta$. Consider a situation where this system is too complicated to solve, but where data are available, usually as sequences of measured values of x , assumed here to be observed with negligible observation errors. One can write $R(x, y)$ in the form

$$R(x, y) = R_0(x) + z(x, y), \quad (2)$$

where R_0 approximates $R(x, y)$ in some sense and is such that one is able to solve the equation

$$\frac{d}{dt}x = R_0(x). \quad (3)$$

The remainder $z(x, y) = R(x, y) - R_0(x)$, often called the “unresolved tendency”, is the contribution of the unresolved variables to the equation for x . Recent work has shown that z can represent model error [8, 9] or model noise [10]; we call z simply the “noise”.

A usual approach to the problem of computing x is to identify z from data [11–13], i.e., find a concise approximate representation \hat{z} of z that can be evaluated on the computer, and then solve the equation

$$\frac{d}{dt}x = R_0(x) + \hat{z}. \quad (4)$$

When \hat{z} is a stochastic process, this is a “stochastic parametrization”. Equation (4) is an instance of a dimensionally-reduced equation, in the sense that it has fewer variables than equations (1). However, this approach has some difficulties. In general the available data are measurements of x , not of z ; to find z so that it can be identified one has to use equation (2), and in particular differentiate x numerically, which is generally impractical or inaccurate because z may have high-frequency components or fail to be sufficiently smooth, and the data may not be available at sufficiently small time intervals (an illuminating analysis in a special case can be found in [14, 15]). If one can successfully identify z , equation (4) generally becomes a nonlinear stochastic differential system, where in general \hat{z} at a given time depends on earlier values of x and \hat{z} (see the next section), that may be hard to solve with sufficient accuracy (see e.g. [16, 17]).

Here we take a different approach. We note that equations (3) and (4) are always solved on the computer, i.e., in discrete form, that the data are always given at a discrete collection of points, and that one wishes to determine x but in general one is not interested in determining z . We can therefore avoid the difficult detour through a continuous z followed by a discretization, by working entirely in the discrete setting as follows. We can pick once and for all a particular accurate discretization of equation (3) with a particular time step δ ,

$$x^{n+1} = x^n + \delta R_\delta(x^n),$$

where R_δ is obtained, for example, from by a Runge–Kutta method, and where n indexes the result after n steps. As in the continuous case, this equation has to be corrected to account for the impact of the continuous variables, and here also for the possible inaccuracy of the numerical scheme. We use the data to identify the discrepancy sequence, $z_\delta^{n+1} = (x^{n+1} - x^n)/\delta - R_\delta(x^n)$, which are available from x data without approximation. This is equivalent to identifying the following reduced system

$$x^{n+1} = x^n + \delta R_\delta(x^n) + \delta z_\delta^{n+1}, \quad (5)$$

which constitutes a discrete stochastic parametrization.

We assume, as one does in the continuous case, that the system under consideration is ergodic, so that its long-time statistics are stationary. The sequence z_δ^n becomes a stationary time series, which we represent by the widely used NARMAX (nonlinear auto-regression moving average with exogenous inputs) representation, with x as an exogenous input. This representation makes it possible to integrate the numerical scheme into the reduced equations, and to take into account efficiently the non-Markovian features of the reduced system as well as model and numerical errors. There is no stochastic differential system to solve. It is important to note that identifying z_δ can be very different from identifying the continuous z . The question, in what sense does z_δ approximate z , is not relevant, since the goal is to calculate x and z_δ is sufficient for the purpose. Note that z_δ should be a good approximation of z whenever equation (3) can be effectively approximated by the first-order Euler scheme. For practical purposes, the discrete stochastic parametrization accomplishes everything that would be accomplished by a successful continuous parametrization followed by an accurate approximation. From now on, we drop the hat that distinguishes between a process and its identification and the subscript δ in z_δ .

The rest of the paper is organized as follows. In the next section we summarize the NARMAX methodology we use, make comparisons with earlier work, and explain the connections with the

Mori-Zwanzig formalism of statistical physics. In section 3 we present our test problem, the Lorenz 96 equations. In section 4 we present numerical results. In the final section we draw conclusions.

2 The NARMAX representation

We represent z in the reduced system (5) by a variant of the NARMAX representation, see e.g. [18–21]. Its advantages are its generality, and the extensive experience in its use. Most of the earlier identifications can be viewed as special case of this representation, as explained below. The greater generality of the NARMAX approach will be increasingly important as model reduction methods are applied to more complex problems. Equation (5) becomes

$$x^{n+1} = x^n + \delta R_\delta(x^n) + \delta z^{n+1}, \quad z^{n+1} = \Phi^{n+1} + \xi^{n+1}, \quad (6)$$

for $n = 1, 2, \dots$, where the ξ^{n+1} are independent Gaussian random variables with mean zero and variance σ^2 , and Φ^n is the sum:

$$\Phi^n = \mu + \sum_{j=1}^p a_j z^{n-j} + \sum_{j=1}^r \sum_{i=1}^s b_{i,j} P_i(x^{n-j}) + \sum_{j=1}^q c_j \xi^{n-j}, \quad (7)$$

μ , σ^2 and $\{a_j, b_{i,j}, c_j\}$ are parameters to be inferred from data, and the exogenous inputs P_i , $i = 1, \dots, s$ are functions to be determined; to simplify the notations, equation (7) has been written as if equation (6) were scalar. This is the NARMAX representation of x and z . In equation (7), the terms in z are the autoregression part of order p , the terms in ξ are the moving average part of order q , and the $P_i(x)$ are the exogenous inputs. Note that in the reduced system (6), the z^n term can be eliminated; indeed, if we substitute $z^n = (x^n - x^{n-1})/\delta - R_\delta(x^{n-1})$ into (7) and the second of equation in (6), we obtain a NARMA representation for x . A suitable choice of the functions P_i will be discussed in the context of a specific example and will connect the representation to the approximation of equation (3).

To implement the NARMAX representation, one has to determine its structure and estimate the parameters. While NARMAX has been widely studied (see e.g. [18, 22–25] and references therein), one should use the earlier work with caution, especially in the detection of structure by least-squares-based methods, because in the standard NARMAX, unlike here, the exogenous process is independent of the noise process. Suppose one has selected a structure, that is, chosen the functions P_i and the orders (p, r, s, q) in the representation (7). Since the representation is linear in the parameters $\theta = (\mu, a_j, b_{i,j}, c_j)$, these parameters can be estimated using conditional likelihoods as follows. The sequence $\{z^n\}$ for $n = 1, 2, \dots, N$ can be computed from the data using the first equation in (6). Once the values of ξ^1, \dots, ξ^q are known, the noise sequence ξ^n for $q+1 \leq n \leq N$ can be computed from $\xi^n = z^n - \Phi^n$. This leads to the conditional log-likelihood of the observations x^n for $q+1 \leq n \leq N$ (up to a constant):

$$l(\theta|\xi^1, \dots, \xi^q) = - \sum_{n=q+1}^N \frac{|z^n - \Phi^n|^2}{2\sigma^2} - \frac{N-q}{2} \ln \sigma^2.$$

Since the NARMAX representation is linear in its parameters other than σ^2 , the log-likelihood $l(\theta|\xi^1, \dots, \xi^q)$ is quadratic in these parameters, its gradient can be easily computed and the maximum likelihood estimator (MLE) can be obtained by standard gradient-based optimization methods, such as the quasi-Newton method. If the reduced system is ergodic, the MLE is asymptotically

consistent, and the initial values of ξ^1, \dots, ξ^q do not affect the result (see e.g. [21, 24]). For convenience, we set $\xi^1 = \dots = \xi^q = E[\xi^1] = 0$.

We remark that for the above Gaussian likelihood, the MLE is the same as the least squares estimator for the parameters, which has been widely used in control (see e.g. [22, 23]). As shown in [22], the above estimation procedure can be made recursive; also, the noise sequence need not be Gaussian.

The detection of the representation's structure, however, is less straightforward, as is well-known, see [18, 20]. Because in our problem the exogenous processes are not independent of the noise, popular techniques such as orthogonal least squares and error reduction ratios (see e.g. [18] and references therein), do not work. We use the following criteria for selecting the structure of the representation: (1) it should fit the long-term statistics of the resolved variables, such as the mean and autocorrelation function; (2) as the size of the data increase, the estimated parameters should converge; (3) the estimated parameters should reflect features of the resolved variables, such as symmetry properties. We find structures that satisfy these criteria by trial and error.

It is of interest to relate the NARMAX representation to the Mori-Zwanzig (MZ) formalism [4, 5, 26, 27]. The overall setting in the MZ representation is the same as here: one has data α for the x variables, and one samples data β for y from a given initial pdf. The MZ formalism creates the approximation $R_0(x)$ in equation (3) by conditional expectation:

$$R_0(x) = E[R(x, y)|x],$$

where $E[a|b]$ is the expected value of a with respect to the initial measure given b . When the system is ergodic and the initial pdf for β is the invariant measure conditioned by α , this is the best least-squares approximation of $R(x, y)$ by a function of x . The MZ formalism then yields an expression for $z(x, y)$ in equation (2) as a sum of a noise and a non-Markovian memory/dissipation term, corresponding to ξ^{n+1} and Φ^{n+1} in equation (7); note that in (7) R_0 is not restricted to the MZ recipe. The MZ expressions are exact, and prove the need for the representation of z to take the memory into account by including information from earlier steps.

Once the initial data $y(0) = \beta$ have been sampled, the MZ equations are deterministic; the MZ formalism proposes to follow in time one particular sample path of the system. For a chaotic system such as the one discussed in the next section, this may not be computationally feasible. The representation here looks for sample paths of a stationary stochastic process whose statistics agree with the statistics defined by the equations of motion. This is a related but less ambitious and more feasible task.

The evaluation of the various terms in the MZ formalism is difficult; as far as we know, there is only one case in the literature [28] where it has been successfully carried out in full for a nonlinear problem that is not completely trivial. The MZ formalism is a useful starting point for analytic approximations (see e.g. [28, 29]) but it is difficult to use it to construct reduced models from data when suitable analytic approximations are not available. The formalism here is a more tractable way to use data for dimensional reduction, and generalizes the MZ formalism to a broader class of approximations.

There is a large literature on data-based dimensional reduction. In [12], see also [30], the noise z is represented as the sum of an approximating polynomial in x obtained by regression and a one-step autoregression; the details are in the next section where we compare our results to those in [12]. The shortcomings of this representation as a general tool are that it does not allow z to remember past values of x , and that the autoregression term is not necessarily small, making it difficult to solve the equations accurately. Furthermore, in [12], numerical values of a continuous z are obtained by finite differences. In [13, 31] the noise is represented by a conditional Markov

chain that depends on both current and past values of x ; the Markov chain is deduced from data by binning and counting, assuming that exact observations of z are available. It should be noted that the Markov chain representation is intrinsically discrete, making this work close to ours in spirit. In [32] the noise is treated as continuous and represented in a form that is partly analogous to the NARMAX representation once one translates from the continuum to the grid. An earlier construction of a reduced approximation can be found in [10], where the approach was not yet fully discrete. Other interesting related work can be found in [33–36].

3 The Lorenz 96 equations

The Lorenz 96 equations [37] are a set of chaotic differential equations that is often used as a metaphor for the atmosphere. It has been widely used as a test bench for various dimensional reduction and stochastic parametrization methods [12, 13, 30, 31, 38]. Following [13], we use the following formulation of the equations:

$$\begin{aligned}\frac{d}{dt}x_k &= x_{k-1}(x_{k+1} - x_{k-2}) - x_k + F + z_k, \\ \frac{d}{dt}y_{j,k} &= \frac{1}{\varepsilon}[y_{j+1,k}(y_{j-1,k} - y_{j+2,k}) - y_{j,k} + h_y x_k]\end{aligned}$$

with $z_k = \frac{h_x}{J} \sum_j y_{j,k}$, and $k = 1, \dots, K$, $j = 1, \dots, J$. The indices are cyclic, $x_k = x_{k+K}$, $y_{j,k} = y_{j,k+K}$ and $y_{j+J,k} = y_{j,k+1}$. The system is invariant under spatial translations, and the statistical properties are identical for all x_k . The formulation here is equivalent to the original formulation by Lorenz (see e.g. [31, 38]); the parameter ε measures the time-scale separation between the resolved variables x_k and the unresolved variables $y_{j,k}$. We set $\varepsilon = 0.5$, so that there is no significant time scale separation between resolved and unresolved processes, as is both more realistic and more difficult to handle for parameterizations (see [38] and references therein). The other parameters are $K = 18$, $J = 20$, $F = 10$, $h_x = -1$ and $h_y = 1$. The ergodicity of the Lorenz 96 system has been numerically verified in earlier work (see e.g. [38]) and we do not revisit this issue.

In the following section we present numerical results produced by our NARMAX scheme and compare them to those in [12] labeled POLYAR. We do not compare with the results of [13] because they require exact observations of z . In [12], the continuous z is estimated by finite differences:

$$z_k(t) \approx \frac{x_k(t + \delta) - x_k(t)}{\delta} - x_{k-1}(x_{k+1} - x_{k-2}) + x_k - F.$$

Then a polynomial regression of the form $z_k(t) = P(x_k(t)) + \eta_k(t)$ is used to fit the data $\{x_k(n\delta), z_k(n\delta)\}$, where $P(x)$ is an approximating polynomial obtained by least squares, and $\eta_k(t)$ is a one-step autoregression (AR(1)) with parameters estimated from the time series $z_k(n\delta) - P(x_k(n\delta))$, for $n = 1, 2, \dots$. This leads to the following reduced stochastic equation:

$$\frac{d}{dt}x_k = x_{k-1}(x_{k+1} - x_{k-2}) - x_k + F + P(x_k) + \eta_k, \quad (8)$$

where η_k is an autoregression of the form

$$\eta_k(t + \delta) = \phi \eta_k(t) + \sigma \xi_k(t) \quad (9)$$

where ϕ, σ are constants deduced from the data, and the $\xi_k(t)$ are independent identically distributed Gaussian random variable with mean zero and variance one, for each component $k = 1, \dots, K$ of the equation. This reduced system is solved as follows: given the current time vectors

Table 1: The estimated parameters in the POLYAR system; the column labeled j for $j = 0, \dots, 5$ contains the coefficient of x to the power j .

	5	4	3	2	1	0	ϕ	σ
$\delta = 0.01$	-0.00002	0.0004	-0.0002	-0.0258	-0.3567	0.0529	0.9948	0.9397
$\delta = 0.05$	-0.00003	0.0009	-0.0035	-0.0137	-1.0030	1.3919	0.7626	11.3857

Table 2: The estimated parameters in the NARMAX model.

	a_1	$b_{1,1}$	$b_{1,2}$	d_1		μ	σ^2
$\delta = 0.01$	0.9782	-0.1271	0.1132	0.9997		0.0115	0.0004
	a_1	$b_{1,1}$	$b_{1,2}$	$b_{1,3}$	$c_{1,1}$	μ	σ^2
$\delta = 0.05$	0.8879	-0.0712	-0.0002	0.0002	-0.0084	0.0556	0.0284

$(\eta_k(t), x_k(t))$, the next time-step $\eta_k(t + \delta)$ is calculated from (9), and then $x_k(t + \delta)$ is computed by integrating (8) by a fourth-order Runge-Kutta methods, with $\eta(t)$ kept constant during each time step.

In the NARMAX scheme, we use the representation (7), choosing one of the functions $P_i(x)$ to be $R_\delta(x)$ from the approximation (3) and the others to be powers of x . This connects the numerical scheme with the representation of the noise. We select the structure and estimate the parameters as described earlier. The parameters are the same for each spatial component, reflecting the spatial symmetry of the equation. Each component of z remembers only its own past and the past of the corresponding component of x . The term Φ^n in equation (7) becomes:

$$\begin{aligned} \Phi^n = & \mu + \sum_{j=1}^p a_j z^{n-j} + \sum_{j=1}^r \sum_{l=1}^{d_x} b_{j,l} (x^{n-j})^l \\ & + \sum_{j=1}^s \sum_{l=1}^{d_R} c_{j,l} (R_\delta(x^{n-j}))^l + \sum_{j=1}^q d_j \xi^{n-j}. \end{aligned} \quad (10)$$

The determination of the numerical parameters in this representation is part of the calculation and the values used are listed in the next section.

4 Numerical results

In the numerical runs, we generate data for x_k by integrating the full Lorenz 96 system with parameters $(\varepsilon, K, J, F, h_x, h_y) = (0.5, 18, 20, 10, -1, 1)$, using a fourth order Runge-Kutta method with time step 0.001. We consider two cases: one in which the observations are made at intervals of $\delta = 0.01$ and one at which they are made at intervals $\delta = 0.05$; the first case corresponds to a case discussed in ([13, 31]); in the second case, the data are slightly sparser. In each case, we make observations at $N = 5 \times 10^5$ points; this requires that the full system be integrated for 5000 and 25000 time units, respectively.

For the POLYAR equations of [12], a fifth-order polynomial is used for both observation settings. Increasing the order further produces small coefficients for the higher degree terms, which do not reduce the variance of noise. The estimated parameters, i.e. the coefficients of the polynomial and the parameters for the autoregression, for the first component x_1 of x are shown in Table 1.

For the NARMAX equation (10), we estimated $(p, r, s, q) = (1, 2, 0, 1)$ and $(d_x, d_R) = (1, 0)$ for

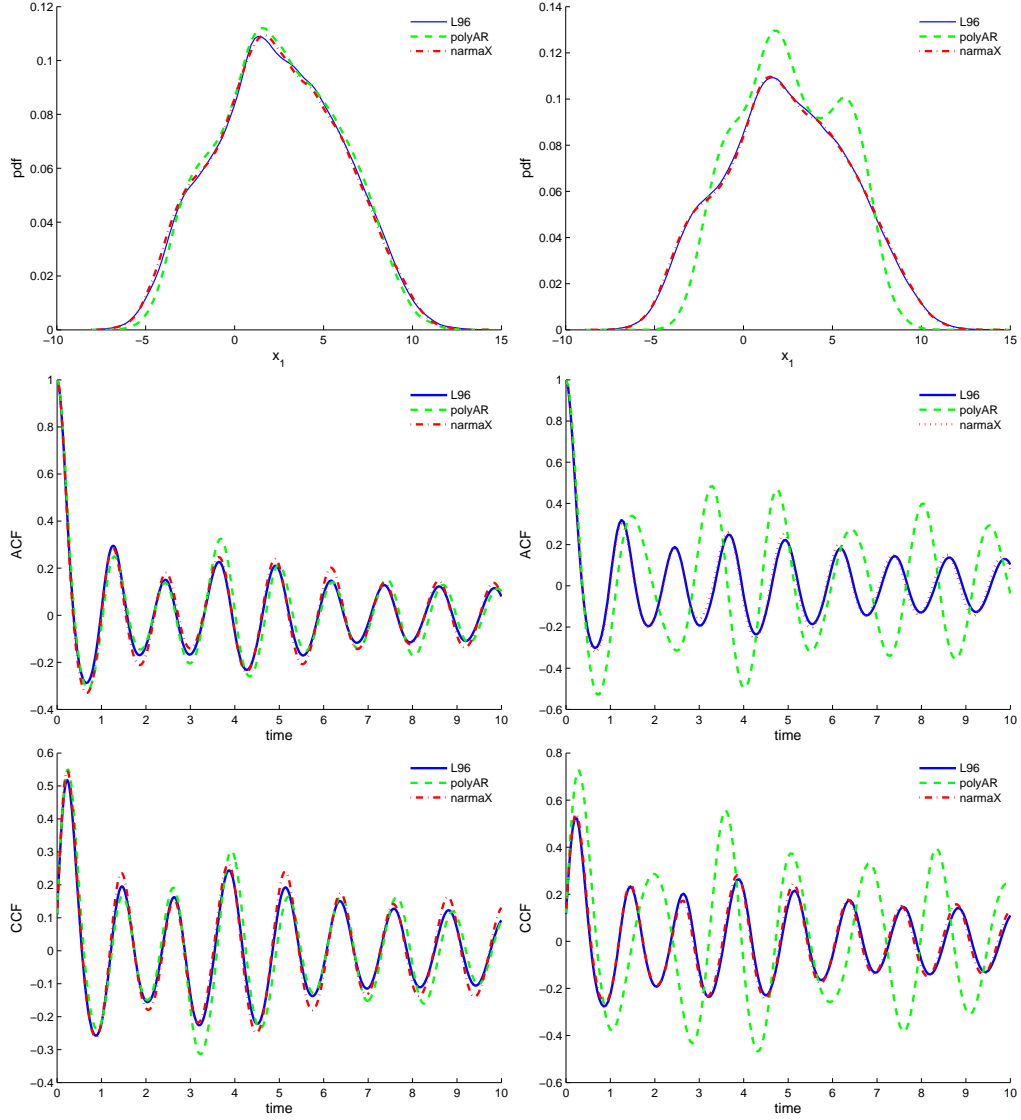


Figure 1: Probability density functions (pdf), autocorrelation functions (ACF) and cross-correlation functions (CCF) of x_1 , produced by the full Lorenz 96 model, the POLYAR system and the NARMAX system. The left column is the case $\delta = 0.01$, and the right column is the case $\delta = 0.05$.

Table 3: The mean, standard deviation and the Kolmogorov–Smirnov statistic (D).

	$\delta = 0.01$			$\delta = 0.05$		
	mean	std.	D	mean	std.	D
full Lorenz 96	2.4506	3.5230		2.3978	3.5222	
POLYAR	2.5335	3.3807	0.0183	2.6031	2.8564	0.0747
NARMAX	2.4113	3.5270	0.0055	2.4293	3.5402	0.0049

the case $\delta = 0.01$, and for the case $\delta = 0.05$, $(p, r, s, q) = (1, 1, 1, 0)$ and $(d_x, d_R) = (3, 1)$. The estimated parameters for x_1 are shown in Table 2.

First, we compare the statistics of the solutions generated by the two reduced systems with the statistics of the full system. We integrate the reduced equations in both cases and obtain values at 5×10^5 points. We calculate the following quantities from the reduced equations as well as from the full Lorenz 96 equations: the mean, the standard deviation, the probability density function (pdf), the Kolmogorov–Smirnov statistic, the autocorrelation function (ACF) of x_1 , and the cross-correlation function (CCF) between x_1 and x_2 . The pdf of x_1 for the full Lorenz 96 system is well reproduced by both reduced systems when $\delta = 0.01$, see left top in Figure 1. In the sparser data case $\delta = 0.05$ the NARMAX equations produce a much better pdf than the POLYAR equations, see right top in Figure 1. Table 3 displays the mean, the standard deviation, and the Kolmogorov–Smirnov statistic that compare the cumulative distributions of the full Lorenz 96 system and with that of the reduced equations. The NARMAX system is more accurate than the POLYAR system in both cases. The autocorrelation function and the cross correlation function are well reproduced by both reduced systems when $\delta = 0.01$, see left middle and left bottom of Figure 1. When $\delta = 0.05$, however, the autocorrelations and the cross correlations reproduced by the POLYAR miss the amplitude of oscillation and exhibit a phase shift from those of the full Lorenz 96 equations, while the NARMAX system remains accurate, see right middle and right bottom of Figure 1.

We now investigate how well a reduced system predicts the behavior of the full system by calculating mean trajectories of the reduced systems and comparing them with a true trajectory of the full Lorenz 96 system, as follows. First we integrate the full Lorenz 96 system for $10 \times N_0$ time units, and store the results as N_0 short trajectories of 10 time units each. For each short true trajectory, we perform N_{ens} integrations of the reduced systems over 10 time units, starting all ensemble members from the same several-step initial conditions as the corresponding full solution – several initial steps are needed to initialize η in POLYAR and ξ in NARMAX. We do not introduce artificial perturbations into the initial conditions, because the exact initial conditions for x are known, and by initializing η and ξ from data, we preserve the memory of the system so as to generate better ensemble trajectories. We then average the solutions of the reduced equations in each subinterval and compare these averages with the trajectories of the full system by calculating the root-mean-square error (RMSE) and anomaly correlation (ANCR) between them; the former measures the average difference between trajectories while the latter measures the average correlation between them; the formulas and their rationale can be found e.g. in [13].

Results for RMSE and ANCR are shown in Figure 2, where we use $N_0 = 10000$, $N_{ens} = 1, 5, 20$ and the number of steps where initial conditions are imposed is $n_0 = \max\{1, p, r, s, 2q\} + 1$. In the case $\delta = 0.01$, the NARMAX reduction performs slightly better than the POLYAR reduction. In the case $\delta = 0.05$, the NARMAX reduction provides a significant improvement over the POLYAR reduction. For example, the forecast lead time at which the anomaly correlation drops below 0.6 is extended from $\tau = 2$ to $\tau = 4$ in the case $N_{ens} = 20$.




Figure 2: Root mean square error (RMSE) and anomaly correlation (ANCR) of ensemble forecasting, produced by the NARMAX system (solid line) and the POLYAR system (dashed line), for different ensemble sizes: $N_{ens} = 1$ (circle marker), $N_{ens} = 5$ (cross marker), and $N_{ens} = 20$ (asterisk marker). The left column is the case $\delta = 0.01$, and the right column is the case $\delta = 0.05$.

5 Conclusions

We have presented a discrete approach to data-based dimensional reduction and stochastic parametrization, in which the problem is consistently treated as discrete, obviating earlier difficulties in estimating noise from measurements and in approximating reduced continuum equations. Within this discrete approach, we have identified the reduced system via the NARMAX representation. This generalizes earlier work, in particular by making it easy to include memory effects in full. We have tested the resulting algorithm on the Lorenz 96 system of equations, often used as a test bench for dimensional reduction schemes; our construction did better than earlier work in reproducing the dynamics and the long-range statistics of the variables of interest, most conspicuously in a problem where the data were sparse. We related our representation to the Mori-Zwanzig formalism and suggested that our methods can be used to construct data-based implementations of this formalism. We expect the advantages of our modeling to become even more marked as it is applied to increasingly complex problems.

6 Acknowledgements

This work was supported in part by the Director, Office of Science, Computational and Technology Research, U.S. Department of Energy, under Contract No. DE-AC02-05CH11231, and by the National Science Foundation under grants DMS-1217065 and DMS-1419044. The authors would like to thank their Berkeley colleague Dr. Matthias Morzfeld, Prof. Kevin Lin of the University of Arizona, Prof. Xuemin Tu of the University of Kansas, and Prof. Robert Miller of the State University of Oregon for helpful comments and good advice.

References

- [1] E. Kalnay. *Atmospheric Modeling, Data Assimilation, and Predictability*. Cambridge University Press, Cambridge, 2003.
- [2] D.S. Wilks. *Statistical Methods in the Atmospheric Sciences*. Academic, Oxford UK, 2011.
- [3] Y. Zeng and S. Wu (eds.). *State Space Models, with Applications in Economics and Finance*. Springer, New York, 2013.
- [4] D. Evans and G. Morris. *Statistical Mechanics of Nonequilibrium Liquids*. Academic Press, London, 1990.
- [5] A.J. Chorin and O.H. Hald. *Stochastic Tools in Mathematics and Science*. Springer, 3rd edition, 2013.
- [6] P. Bernard and J. Wallace. *Turbulent Flow: Analysis, Measurement, and Prediction*. Wiley, Hoboken NJ, 2002.
- [7] A.J. Chorin. *Vorticity and Turbulence*. Springer, NY, 1994.
- [8] J. Harlim. Data assimilation with model error from unresolved scales. *arXiv:1311.3579*, 2013.
- [9] T. Berry and J. Harlim. Linear theory for filtering nonlinear multiscale systems with model error. *Proc. R. Soc. A*, 470:20140168, 2014.
- [10] A.J. Chorin and O.H. Hald. Estimating the uncertainty in underresolved nonlinear dynamics. *Math. Mech. Solids*, 19:28–38, 2014.
- [11] G.A. Pavliotis and A.M. Stuart. Parameter estimation for multiscale diffusions. *J. Statist. Phys.*, 127(4):741–781, 2007.

- [12] D.S. Wilks. Effects of stochastic parameterizations in the Lorenz'96 model. *Q. J. R. Meteorol. Soc.*, 131:389–407, 2005.
- [13] D. Crommelin and E. Vanden-Eijnden. Subgrid-scale parameterization with conditional Markov chains. *J. Atmos. Sci.*, 65:2661–2675, 2008.
- [14] Y. Pokern, A.M. Stuart, and P. Wiberg. Parameter estimation for partially observed hypoelliptic diffusions. *J. Roy. Statist. Soc. B*, 71(1):49–73, 2009.
- [15] A. Samson and M. Thieullen. A contrast estimator for completely or partially observed hypoelliptic diffusion. *Stochastic Process. Appl.*, 122(7):2521–2552, 2012.
- [16] P.E. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*. Springer, Berlin, 3rd edition, 1999.
- [17] G.N. Milstein and M.V. Tretyakov. *Stochastic numerics for mathematical physics*. Springer-Verlag, 2004.
- [18] S.A. Billings. *Nonlinear system identification: NARMAX methods in the time, frequency, and spatiotemporal domains*. John Wiley and Sons, 2013.
- [19] P. Brockwell and R. Davis. *Introduction to Time Series and Forecasting*. Springer, New York, 2002.
- [20] J. Fan and Q. Yao. *Nonlinear time series: nonparametric and parametric methods*. Springer, New York, 2003.
- [21] J.D. Hamilton. *Time Series Analysis*. Princeton University Press, 1994.
- [22] F. Ding and T. Chen. Identification of Hammerstein nonlinear ARMAX systems. *Automatica*, 41(9):1479–1489, 2005.
- [23] H.F. Chen. New approach to recursive identification for ARMAX systems. *IEEE Trans. Automat. Contr.*, 55(4):868–879, 2010.
- [24] E.J. Hannan. The identification and parameterization of ARMAX and state space forms. *Econometrica*, pages 713–723, 1976.
- [25] D.S. Stoffer. Estimation and identification of space-time ARMAX models in the presence of missing data. *J. Am. Statist. Assoc.*, 81:762–772, 1986.
- [26] R. Zwanzig. Nonlinear generalized Langevin equations. *J. Stat. Phys.*, pages 215–220, 1973.
- [27] R. Zwanzig. *Nonequilibrium Statistical Mechanics*. Oxford, 2001.
- [28] A.J. Chorin, O.H. Hald, and R. Kupferman. Optimal prediction with memory. *Physica D*, 166:239–257, 2002.
- [29] P. Stinis. Renormalized reduced models for singular PDE's. *Comm. Appl. Math. Comp. Phys.*, 8:39–66, 2013.
- [30] H.M. Arnold, I.M. Moroz, and T.N. Palmer. Stochastic parametrization and model uncertainty in the Lorenz'96 system. *Phil. Trans. R. Soc. A*, 371:20110479, 2013.
- [31] F. Kwasniok. Data-based stochastic subgrid-scale parametrization: an approach using cluster-weighted modeling. *Phil. Trans. R. Soc. A*, 370:1061–1086, 2012.
- [32] A.J. Majda and J. Harlim. Physics constrained nonlinear regression models for time series. *Nonlinearity*, 26(1):201–217, 2013.
- [33] M.D. Chekroun, D. Kondrashov, and M. Ghil. Predicting stochastic systems by noise sampling, and application to the El Niño-Southern Oscillation. *Proc. Natl. Acad. Sci. USA*, 108:11766–11771, 2011.
- [34] D. Kondrashov, M.D. Chekroun, and M. Ghil. Data-driven non-Markovian closure models. *Physica D: Nonlinear Phenomena*, 297:33–55, 2015.

- [35] J. Duan and B. Nadiga. Stochastic parameterization for large eddy simulation of geophysical flows. *Proc. Amer. Math. Soc.*, 135(4):1187–1196, 2007.
- [36] A. Du and J. Duan. A stochastic approach for parameterizing unresolved scales in a system with memory. *J Algorithm Comput Technol.*, 3(3):393–405, 2009.
- [37] E.N. Lorenz. Predictability: A problem partly solved. *Proc. Seminar on predictability*, pages 1–18, 1995.
- [38] I. Fatkulin and E. Vanden-Eijnden. A computational strategy for multi-scale systems with applications to Lorenz’96 model. *J. Comput. Phys.*, 200:605–638, 2004.